



Accelerator Science and Technology

□ From SLAC Web Site ...

SLAC Experiment Identifies New Subatomic Particle

Physicist Antimo Palano representing the BABAR experiment presented the evidence for the identification of a new subatomic particle named D_s (2317) to a packed auditorium on Monday, April 28 at SLAC. Initial studies indicate that the particle is an unusual configuration of a "charm" quark and a "strange" anti-quark.





Accelerator Science and Technology

- Accelerators are important.
 - Research in particle physics.
 - Fundamental to understanding of structure of matter.



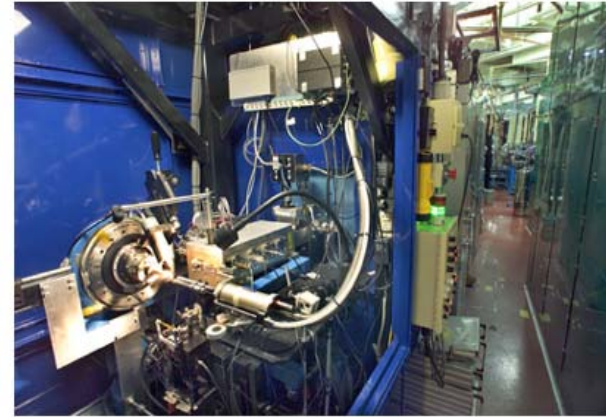
- Accelerators are expensive.
 - High cost in construction, operations, and maintenance.
 - Represent major DOE investment.



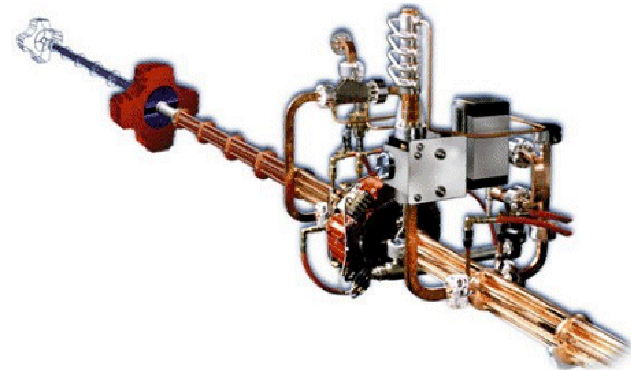


Accelerator Science and Technology

- ❑ Accelerator simulations and modeling are indispensable.
 - Understanding the science of accelerators for safe operations.
 - Improving performance and reliability of existing accelerators.
 - Designing next generation of accelerators accurately and optimally.



ALS Beamline 8.3.1
VERSION 1.1.0 07/05/00



Terascale Accelerator Modeling

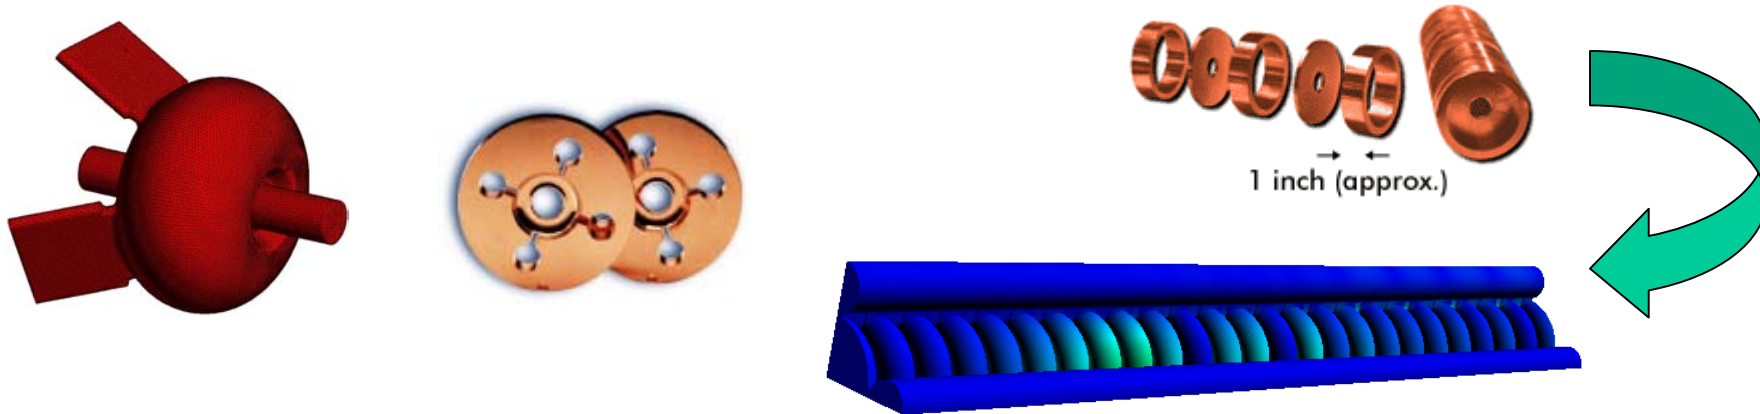


- ❑ Three components in the SciDAC Project on Accelerator Science and Technology.
 - Beam Systems Simulations (R. Ryne, LBNL).
 - Electromagnetic Systems Simulations (K. Ko, SLAC).
 - Advanced Accelerator Systems Simulations (W. Mori, UCLA).

- ❑ TOPS is actively collaborating with Electromagnetic Systems Simulations at SLAC.
 - Linear Algebra - large-scale sparse eigensolvers, sparse linear equations solvers (LBNL, Stanford, SLAC).
 - Load Balancing - improving performance and scalability (LBNL, SLAC, Sandia).



Designing Accelerator Structures



- Modeling of accelerator structures requires the solution of the Maxwell equations.
 - Finite element discretization in frequency domain leads to a large sparse generalized eigenvalue problem.

$$K x = \lambda M x, \quad K \geq 0; M > 0$$



Designing Accelerator Structures

- ❑ Design of accelerator structures.
 - Modeling of a single accelerator cell suffices.
 - Relatively small eigenvalue problem.
 - There is an optimization problem here ...
 - But need fast and reliable eigensolvers at every iteration.
- ❑ Understanding the wake field requires the modeling of the full structure.
 - Need to compute a large number of frequency modes.

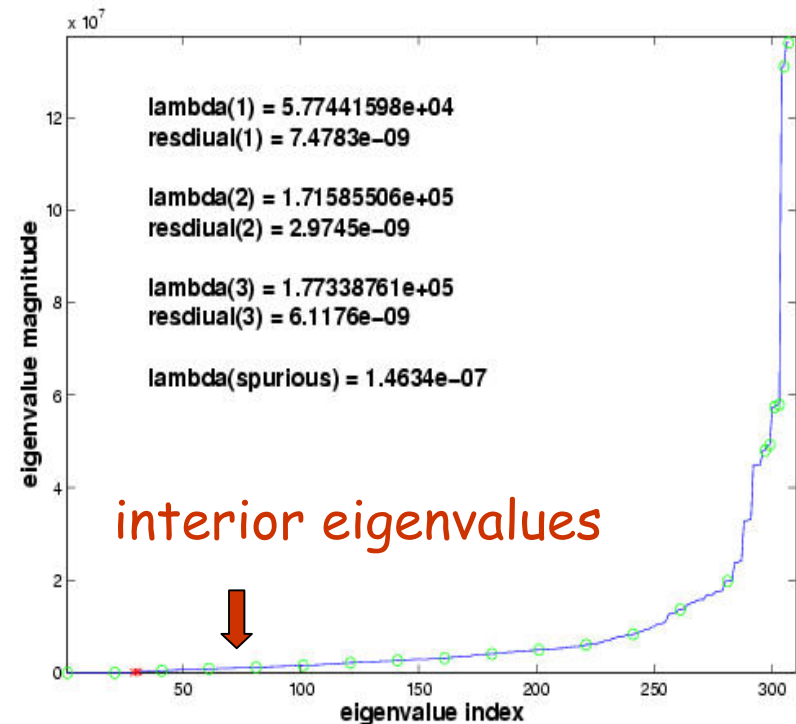




Challenges in Eigenvalue Calculations

- ❑ 3-D structures \Rightarrow large matrices.
- ❑ Need very accurate interior eigenvalues that have relatively small magnitudes.
- ❑ Eigenvalues are tightly clustered.
- ❑ When losses in structures are considered, the problems will become complex symmetric.
- ❑ Omega3P has been able to compute eigen modes of a 82-cell structure with 22M DOF's (without losses).

Spectral Distribution





Large-scale Eigenvalue Calculations

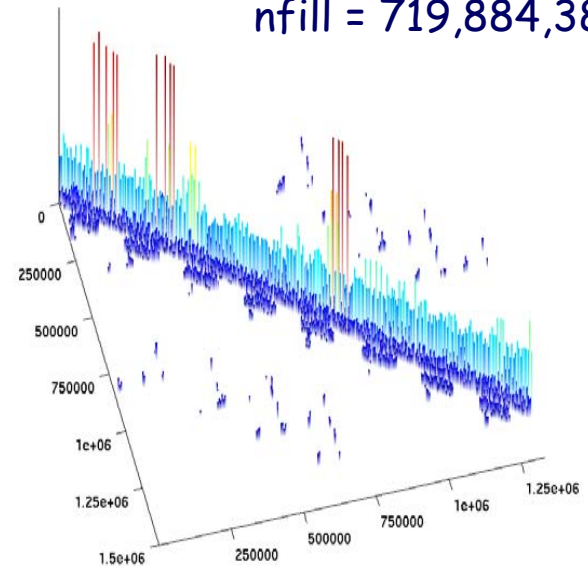
- ❑ Parallel shift-invert Lanczos algorithm.
 - Ideal for computing interior and clustered eigenvalues.
$$K x = \lambda M x \quad \rightarrow \quad M(K - \sigma M)^{-1} M x = \mu M x$$
 - Need solution of sparse linear systems.
 - SLAC: inexact solution + Newton-type correction (Omega3P).
- ❑ Exact shift-invert Lanczos - require complete factorizations of (sparse) matrices.
 - Make possible by exploiting work on sparse direct solvers in TOPS.
 - Combine SuperLU_DIST with PARPACK to obtain a parallel implementation of a shift-invert Lanczos eigensolver.
 - Enable accurate calculation of eigenvalues, allow verification of other eigensolvers, and provide a baseline for comparisons.

TOPS Contribution - SuperLU



- ❑ SuperLU and SuperLU_Dist.
 - Direct solution of sparse linear system $Ax = b$.
 - Efficient, high-performance, portable implementations on modern computer architectures.
 - Support real and complex matrices, fill-reducing orderings, equilibration, numerical pivoting, condition estimation, iterative refinement, and error bounds.

dds47 matrix:
 $n = 1,323,019$
 $\text{nnz} = 20,127,775$
 $\text{nfill} = 719,884,387$



TOPS Contribution - SuperLU



- ❑ New developments/improvements in SuperLU are motivated by the accelerator application.
 - Accommodate distributed input matrices.
 - Symbolic factorization still sequential but reduction in memory used.
 - Improve triangular solution routine (in progress).
 - Improve management of buffers used for non-blocking operations to make it friendlier to MPI implementations.
 - Use partial inversion to improve parallelism in the substitution process.

Problem	p	Time (ESIL)	Nonzeros in L+U-I	Time (Hybrid)
dds15 linear (14 eigenvalues)	32	4,413.9	867,709,851	7,430.2
dds47 linear (16 eigenvalues)	48	4,859.8	719,884,387	12,477.8



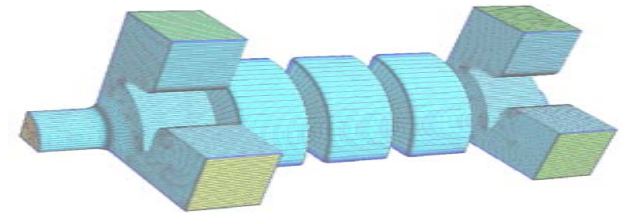
Large-scale Eigenvalue Calculations

- ❑ TOPS' shift-invert Lanczos and Omega3P produce the same eigenvalues.
- ❑ SLAC considers both the exact shift-invert Lanczos and the inexact shift-invert Lanczos as complementary.
- ❑ Exact shift-invert Lanczos is a serious contender because of memory availability on highly parallel machines.
 - Integrated as a run-time option in Omega3P.
- ❑ The exact shift-invert solver provides a quick solution to the sparse complex symmetric eigenvalue problems.

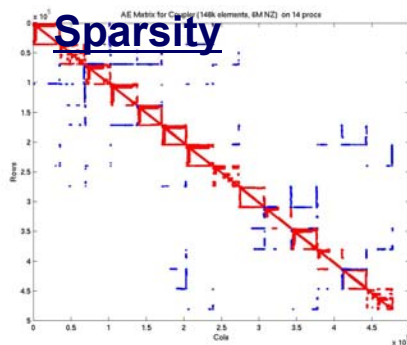


Load Balancing in Time-domain Solver

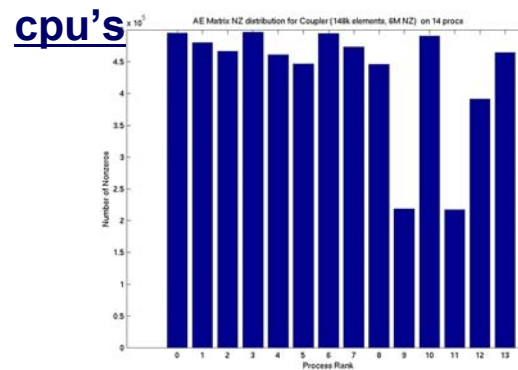
- ❑ Load balancing problem in Tau3P, a time-domain solver.
 - Use of unstructured meshes and refinements lead to matrices for which nonzero entries are not evenly distributed.
 - Makes work assignment and load balancing difficult in a parallel setting.
 - SLAC's Tau3P currently uses ParMETIS to partition the domain to minimize communication.



Matrix

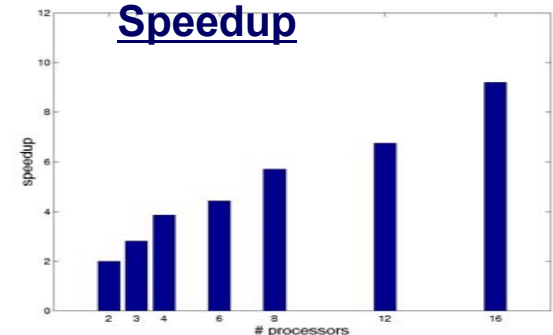


Matrix Distribution over 14



Parallel

Speedup

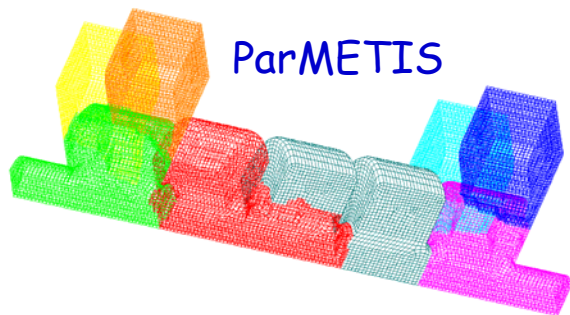




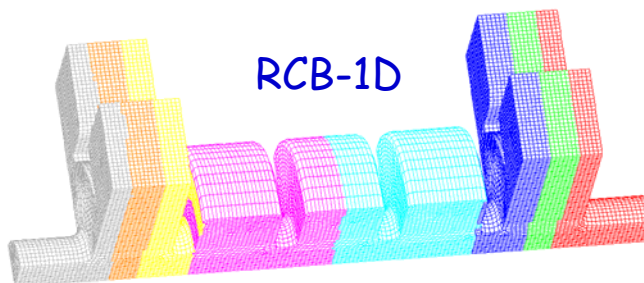
Load Balancing in Time-domain Solver

- ❑ Collaboration between SLAC and TOPS (+ Sandia) has resulted in improved performance in Tau3P.
 - Sandia's Zoltan library is implemented to access better partitioning schemes for improved parallel performance over existing ParMETIS tool through reduced communication costs.

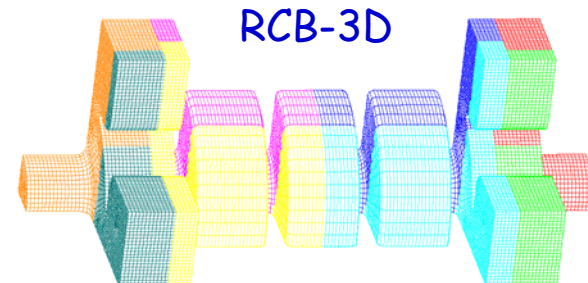
8 processor partitioning of a 5-cell RDDS with couplers on NERSC IBM SP			
	Tau3P Runtime	Max. Adj. Procs.	Max. Bound. Objects
ParMETIS	288.5 sec	3	585
RCB-1D	218.5 sec	2	3128
RCB-3D	345.6 sec	5	1965



ParMETIS



RCB-1D



RCB-3D



Load Balancing in Time-domain Solver

ParMETIS



RCB-1D



Performance results on NERSC IBM SP for a 55-cell structure				
# of processors	ParMETIS run-time	ParMETIS max. adj. procs	RCB-1D run time	RCB-1D max. adj. procs
32	1455.0	4	1236.6	2
64	736.6	4	627.2	2
128	643.0	10	265.1	2
256	360.0	11	129.2	2
512	292.1	14	92.3	4

- ❑ Significant improvement obtained from using RCB-1D over ParMETIS on a 55-cell structure due to the linear nature of the geometry.



Other Activities and Future Plans

- ❑ Sparse direct solvers.
 - More improvements (e.g., symbolic factorization & triangular solutions) to make SuperLU more scalable.
 - Fill-reducing orderings.
 - Scheduling issues.
- ❑ Incomplete factorization algorithms.
 - Exploiting technology from sparse direct methods.
- ❑ Eigenvalue calculations:
 - More comparisons using larger problems in progress.
 - Use of sparse symmetric factorization.
 - Iterative solvers + preconditioning techniques for inexact shift-invert Lanczos.
 - Other eigen solvers (e.g., Jacobi-Davidson, multigrid).
- ❑ Role of optimization techniques.